

The St. Petersburg Paradox

Knut K. Aase
Norwegian School of Economics and
Business Administration
5035 Sandviken - Bergen, Norway

Sept., 1998

Abstract

The classical St. Petersburg Paradox is discussed in terms of doubling strategies. It is claimed that what was originally thought of as a “paradox” can hardly be considered as very surprising today, but viewed in terms of doubling strategies, we get some results that look paradoxical, at least to the practically oriented investor.

KEYWORDS: St. Petersburg Paradox, free lunch, arbitrage possibility, expected utility, certainty equivalent, uniform integrability

1 Introduction

In the early days of the calculus of probability it was taken as granted that the value, and hence the “fair price” of a gamble was the mathematical expectation of the gain. Thus this price would be

$$E\{X\} = \int x dF(x)$$

if X represents the gains of the gamble having cumulative probability distribution function F . Applied to insurance this means that the fair premium p for a risk described by the non-negative random variable X would be

$$p = \int_0^{\infty} x dF(x).$$

The Bernoulli Principle

Daniel Bernoulli (1738) published an example, originally presented to him by his cousin Nicolas Bernoulli, where the above method does not work simply because the integral above does not converge. The example has been known as the *St. Petersburg Paradox*, and deals with a game where a coin is tossed until it shows heads. If the first head appears at the n 'th toss, a prize of 2^n

is paid. The expected gain in this gamble is $+\infty$, and Daniel argued that no rational person would be willing to pay an arbitrary large amount for the right to participate in this gamble. He is in fact more explicit about it, and writes that “there should be no sensible man who would not be willing to sell his right to this gain for 20 ducats”.

As an alternative to the expected gain, Daniel Bernoulli suggested that a person would assign the “moral value” of $\log(x)$ to a gain of x . The value of the gamble was then suggested to be the “moral expectation”

$$E\{\log X\} = \int_0^\infty \log x dF(x) = \log 2 \sum_{n=1}^\infty n \left(\frac{1}{2}\right)^n = 2 \log 2,$$

a finite number. Daniel Bernoulli regarded the “paradox” as resolved, and assigned this finite number as the “price” of the lottery. This is of course a very ad hoc solution, which would e.g., not help if the gain was changed to 2^{2^n} instead of 2^n , which was also realized at the time. This solution is not related to any law of large numbers either. Anders Martin-Löf (1985) has discussed and developed a law of large numbers for the St.Petersburg game, which we will come back to below.

Although the subsequent discussion of the “moral value” of a gain seems rather arbitrary today, this discussion gave the starting point of the expected utility theory, where a more general utility function $u(x)$ replaces $\log x$.

In fact expected utility is not at all used in the manner indicated above, and we find it useful to demonstrate below how the vonNeuman-Morgenstern expected utility theory may be used in the pricing of lotteries, and why Daniel Bernoulli’s solution is wrong. But first we mention some statistics.

The Law of Large Numbers for the St. Petersburg Game

Buffon seems to be the only researcher in the 1700 who really tried to play a large number M of plays and calculate the empirical mean $\frac{S_M}{M} := \frac{1}{M} \sum_{k=1}^M Y_k$, where Y_k is the payoff in play no. k . Here the lottery Y pays a gain of one ducat if heads shows in the first trial, two ducats if first heads appears in the second trial, 4 in the third, 8 in the fourth, etc. He let a child play $2048 = 2^{11}$ rounds and found that the number of plays of length 1, 2, 3, ... was 1061, 494, 232, 137, 56, 29, 25, 8, 6. The expected values of these numbers are 2^{11-k} , $k = 1, 2, 3, \dots$ i.e., 1024, 512, 256, 128, 64, 32, 16, 8, 4, 2, 1, ... The total gain was 10057 yielding an average of $\frac{S_M}{M} = 4.91$. One may wonder if this is close to any fair price of this lottery. Let $M = 2^n$. Then Feller (1968) has shown that $\frac{S_M}{M} / \frac{n}{2} \rightarrow 1$ in probability as $n \rightarrow \infty$, so there is not much hope to find any finite value as the price of the game.

Anders Martin-Löf (1985) has been much more specific, and found a probability distribution $F(x)$ such that

$$P\left(\frac{S_M}{M} \Leftrightarrow \frac{n}{2} \leq x\right) \rightarrow F(x), \quad M = 2^n, n \rightarrow \infty.$$

Furthermore he showed that with good approximation $1 \Leftrightarrow F(x) \approx 2^{-m}$ for $m \geq 5$ so that $P(\frac{SM}{M} > \frac{n}{2} + 2^m) \approx 2^{-m}$. By the help of this approximation he suggests that one could determine a premium per game which has some credibility. If one requires a probability of $10^{-3} \approx 2^{-10}$ the fee should be $\frac{n}{2} + 2^{10} = \frac{n}{2} + 1024$ per game. He also comments that " $\frac{n}{2}$ is small compared to 2^m for reasonable values of n , so that in practice it is possible to determine a premium per game independent of n , just as we are used to for games having finite values". The impressive results of Anders Martin-Löf more or less concludes the probabilistic analysis of this game.

In the next sections we turn to a rather different way of valuing lotteries, and in particular the St. Petersburg game.

2 Certainty Equivalents

In this section we indicate how one may possibly use utility functions to obtain individual values of lotteries. There are other pricing theories which are much more involved, using concepts of equilibrium, but we may in fact get somewhere by simply doing the following: Consider an individual having a wealth w_0 (a real number) and facing a lottery with payoff Y . The individual has a Bernoulli utility function, sometimes called a Bernoulli index, $u : R \rightarrow R$. By this we mean the following: Let \prec be a preference relation on the set of random variables, where \preceq and \sim are derived from \prec in the usual way. If this binary relation satisfies a certain set of axioms, where the independence axiom is the most famous, the preference relation can be shown to have a von Neuman-Morgenstern expected utility representation: $W_1 \preceq W_2 \Leftrightarrow E\{u(W_1)\} \leq E\{u(W_2)\}$ for random wealths W_1 and W_2 .

Let us assume that u is increasing and concave. A certainty equivalent for a lottery Y and initial wealth w_0 is the real number w^* satisfying

$$u(w^*) = E\{u(w_0 + Y)\}.$$

It is natural to define the price (the "bid price") of the lottery by

$$p = w^* \Leftrightarrow w_0. \tag{1}$$

This definition may be motivated from common trade. As the owner of some good the price equals the cash balance after the transaction minus the initial cash balance. Here the good corresponds to the lottery Y and the initial and final cash balances are respectively w_0 and w^* . The above definition thus applies the natural definition of a bid price to a lottery. Here we may emphasize that the bid price p defined above is actually the *minimum* price demanded by the individual to sell the lottery.

Finally let us define the *risk premium* π of the lottery as follows:

$$\pi = EY \Leftrightarrow p.$$

The risk premium tells us how much compensation a risk-averse person requires in order to accept a risk. For a risk-averse person the function u is strictly

concave and the risk premium is positive, while for a risk-lover u is strictly convex and π is negative. A risk-neutral individual has a linear Bernoulli index u , and the corresponding risk premium is zero. The risk neutral case is thus the one referred to at the beginning.

It is obvious that if an individual's preferences over probability distributions can be represented by von Neuman-Morgenstern expected utility with the associated Bernoulli utility function $u(w)$, then an affine transformation $au(w) + b, a > 0, b \in R$ represents the same preferences. A consequence of this should be that the certainty equivalent, and hence the bid price, does not depend upon a or b . This latter fact is easily demonstrated:

Proposition 1 *Consider two individuals with same initial wealth w_0 facing the same lottery Y . Assume one has Bernoulli index $u_1(w)$, the other $u_2(w)$. Then if $u_2(w) = au_1(w) + b$, they assign the same price to the lottery for any $a > 0, b \in R$*

Proof: The bid prices p_1 and p_2 are defined respectively by

$$u_1(w_0 + p_1) = Eu_1(w_0 + Y) \quad (2)$$

and

$$u_2(w_0 + p_2) = Eu_2(w_0 + Y) \quad (3)$$

Using the affine structure of u_2 in equation (3), we get

$$au_1(w_0 + p_2) + b = aEu_1(w_0 + Y) + b,$$

which implies by equation (2) that $u_1(w_0 + p_2) = u_1(w_0 + p_1)$. Since $u_1(w)$ is assumed strictly monotonic, it follows that $p_1 = p_2$. \square

We notice that only the requirement $a \neq 0$ is actually used in the above.

Now we can immediately recognize why Daniel Bernoulli's theory is not in agreement with this use of expected utility. An individual with Bernoulli utility index $u(x) = \log x$ should, according to Proposition 1, assign the same value to the St. Petersburg game as an individual having index $u(x) = 2 \log x + 100$, but in Daniel's theory the first would charge $2 \log 2$, the other $(4 \log 2 + 100)$, etc.

We may now ask what value should be assigned to this lottery according to this principle. Before we attempt an answer, it may be an advantage to take a new look at the St. Petersburg game.

3 The St. Petersburg Paradox as an Arbitrage

Let us here turn to the following interpretation of the St. Petersburg game, suggesting why it can still be considered as a "paradox"¹. Consider an agent using the same "doubling strategy" as above, where the agent pays for the

¹The fact that a random variable X is not a member of L^1 can hardly in itself be considered as a "paradox", where $L^1 = \{X; E\{|X|\} < \infty\}$.

sequence of fair games as he goes along until head appears for the first time. Denote the net gain from the game by X . If e.g., heads appeared for the first time on the third trial, he would by then have paid 1 in the first trial, 2 in the second, 4 in the third, so by the beginning of the third trial he would have paid 7 altogether. If heads then turns up, he is paid $2^3 = 8$, and has hence a net gain of 1, after which he quits the game. The net gain will always be the same, and equal to one, if the game ends with heads, and since the probability that this will happen eventually is equal to one, one seems to have something starting to resemble a real “paradox”. This is indeed an “arbitrage possibility”, sometimes called a “free lunch” in financial terminology.

To see this, consider the state space $\Omega = \{e_1, e_2, \dots\}$, where $e_1 = \text{H}$, $e_2 = \text{TH}$, $e_3 = \text{TTH}$ etc., i.e., $e_n = \{\text{first head happens in the } n\text{'th trial}\}$. Then

$$\text{Probability of eventual success} = \sum_{n=1}^{\infty} P(\{e_n\}) = \sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n = 1.$$

In other words it seems as if playing this game will lead the agent to a certain net gain of 1. This seems puzzling since the sequence of games is fair, so one would believe that the seller of the game would just break even in the long run.

The game can clearly be considered as a stopping problem, where the optimal strategy exists. The problem is it may take a very long time².

Since it may take a long time before heads turns up for the first time, the agent must in reality have an *unbounded fortune* (or unbounded credit).

If Daniel Bernoulli had looked at the game this way, he might have come to the conclusion that the game should cost 1, using the expected value principle, since this also is the net expected gain of the game, i.e., $E\{X\} = 1$. Also note that $P[X = 1] = 1$.

Bid and Ask Prices

Consider a seller (a casino) having a certain wealth w_0 , and Bernoulli utility index $u(x) = \log x$. The casino would face the payoff $Y = \Leftrightarrow X$, where X is the payoff from the St. Petersburg game as explained above. The certainty equivalent w^* for the seller of this game is then computed from

$$\log w^* = E \log(w_0 \Leftrightarrow X) = \sum_{k=1}^{\infty} \log(w_0 \Leftrightarrow 1) \left(\frac{1}{2}\right)^k = \log(w_0 \Leftrightarrow 1),$$

which implies that $w^* = (w_0 \Leftrightarrow 1)$. Thus the (seller's) price p for this lottery is $p = (w^* \Leftrightarrow w_0) = \Leftrightarrow 1$, and the risk premium $\pi = (EY \Leftrightarrow p) = 0$.

The interpretation is as follows: Suppose a casino is *obliged* to offer the game. It is then willing to pay (at most) one unit to someone else to get rid of this obligation.

²If the game continues long enough, time will clearly be a constraint, since each game must be presumed to take at least a certain minimum amount of time to carry through, and no agent has an unlimited time to his disposal.

It can also be interpreted as the price charged from someone, having an infinite fortune or credit limit, to play this game. The risk premium is zero since there is no risk for the seller, so the price is the same as the one obtained under risk neutrality, i.e., the premium that Daniel Bernoulli presumably would have suggested.

A buyer's price p_b of any lottery Y could now be defined as follows:

$$u(w_0) = Eu(w_0 \Leftrightarrow p_b + Y). \quad (4)$$

This price is then the maximal amount a buyer, having a certain fortune w_0 and utility function u , would be willing to pay for the lottery Y . With this entrance fee the buyer is indifferent between his present level of expected utility and the level he obtains after accepting the game at price p_b .

In the present situation the buyer has access to infinite credit, and faces the St. Petersburg game. We find that $p_b = 1$ by a computation similar to the one above³. In this case there is no risk for the buyer to pay the entrance fee of 1 unit, and then start playing. With this fee in place the arbitrage possibility of course disappears.

4 A more realistic version of the St. Petersburg Game: Finite credit

Let us look at the game in more realistic terms, and assume that the agent has a *finite* fortune N at his disposal⁴. For simplicity assume $N := N_m = (2^m \Leftrightarrow 1)$ for some positive integer m . Denote the net gain from this game by X_m . First observe that the sequence of random variables $\{X_m; m \geq 1\}$ converges to X in probability as $m \rightarrow \infty$ (notation: $X_m \xrightarrow{P} X$), and also almost surely (notation: $X_m \xrightarrow{a.s.} X$). Now, for any m

$$E\{X_m\} = 1 \cdot \sum_{n=1}^m P(\{\epsilon_n\}) \Leftrightarrow N_m \sum_{n=m+1}^{\infty} P(\{\epsilon_n\}) = 0.$$

Thus the entrance fee for playing this game should be 0, at least according to the “expected value principle”, we have no longer a free lunch and are back in the real world. Still the agent has a relatively large probability of winning 1 if m is large, but he has the small probability $(\frac{1}{2})^m$ of loosing his entire fortune N_m , a very large quantity if m is large enough.

Let us now see what happens if his fortune N increases beyond any limit. Will we then come back to the “free lunch”- situation described above? Since $E\{X_m\} = 0$ for all m , clearly

$$0 = \lim_{m \rightarrow \infty} E\{X_m\} \neq E\{\lim_{m \rightarrow \infty} X_m\} = E\{X\} = 1,$$

³One may notice that we abstract from the time depreciation of money, since it may take some time before the certain gain of 1 is realized.

⁴Discussions with Frøystein Gjesdal are greatly acknowledged on this issue.

which means that we are not back! This might seem puzzling at first: By starting with a large, but finite fortune, it is not possible to get from the situation with “no free lunch” to the situation with arbitrage possibilities by simply increasing this fortune beyond any limit. One has to start at the outset with this unbounded fortune in order to obtain a “free lunch”.

In mathematical terms we have found a situation where we may not pass the limit inside the expectation: Here the sequence of random variables $\{X_m; m \geq 1\}$ converges to X in probability, but the sequence $\{X_m; m \geq 1\}$ does not converge in L_1 -norm. In other words, the sequence $\{X_m; m \geq 1\}$ cannot be *uniformly integrable*, because if it were, we would have been able to pass the limit inside the expectation above. A mathematician would again not call this a paradox, but rather a neat counterexample. It illustrates that while mathematicians may treat limits and infinity with great ease ⁵, when applied to practical situations one has to be really careful; that is where philosophy enters.

The Bid Price

Let us now apply our pricing theory outlined above to this case. First consider the seller (a casino): Here the lottery $Y = \Leftrightarrow X_m$, and the certainty equivalent w^* satisfies

$$\begin{aligned} \log w^* &= E \log(w_0 \Leftrightarrow X_m) \\ &= \sum_{k=1}^m \log(w_0 \Leftrightarrow 1) \left(\frac{1}{2}\right)^k + \log(w_0 + (2^m \Leftrightarrow 1)) \left(\frac{1}{2}\right)^m. \end{aligned} \tag{5}$$

Thus if the success occurs before the m -th play, the seller has to pay 1 unit to the player, but in the case where the player’s fortune runs out before the first heads appears, the casino keeps his entire fortune N . It follows that

$$w^* = (w_0 \Leftrightarrow 1)^{(1 - (\frac{1}{2})^m)} (w_0 + N)^{(\frac{1}{2})^m}. \tag{6}$$

From this expression and the definition of the bid price in equation (1) we can infer that the price p of the casino is in $(\Leftrightarrow 1, 0)$. This means that the price the casino charges, $\Leftrightarrow p$, is here less than 1, the price in the previous case, since it is a possibility that the casino can net the amount N on the game - if luck runs out for the player. Also the price $\Leftrightarrow p > 0$ simply because of risk aversion, since the utility function of the seller is assumed to be $u(x) = \log x$, a concave function.

The Ask Price

Finally consider the buyer. Again making the same assumptions as before regarding preferences, we must now assume that his certain fortune $w_0 > N + p_b$ in order for the expected utility to be well-defined. His price p_b is determined

⁵in e.g., nonstandard theory

by the equation

$$\begin{aligned} \log w_0 &= E \log(w_0 \Leftrightarrow p_b + X_m) \\ &= \sum_{k=1}^m \log(w_0 \Leftrightarrow p_b + 1) \left(\frac{1}{2}\right)^k + \log(w_0 \Leftrightarrow p_b \Leftrightarrow (2^m \Leftrightarrow 1)) \left(\frac{1}{2}\right)^m. \end{aligned} \quad (7)$$

We find that p_b must satisfy the equation

$$w_0 = (w_0 \Leftrightarrow p_b + 1)^{(1 - (\frac{1}{2})^m)} (w_0 \Leftrightarrow p_b \Leftrightarrow N)^{(\frac{1}{2})^m}. \quad (8)$$

From this we observe that the buyers price p_b is smaller than 0, the price under risk neutrality. A negative value of p_b means that the buyer *must be offered* at least a positive side-payment of $(\Leftrightarrow p_b) > 0$ to play the game, and happens because the expected payoff is not large enough to compensate the risk averse buyer for the risk involved.

Notice that we have not found a market price in this case. Even if the seller is risk-neutral, the buyer would not accept. The buyer must in fact be risk-neutral in order to accept this gamble at the “fair price” of zero, and he must be risk-loving to accept the gamble described above.

5 Concluding Remarks

In daily life some firms (investment banks or other financial institutions) seem to routinely play this game from time to time. On a few occasions the results of such games also make the headlines of newspapers around the world. These firms, or the dealers who trade on behalf of the firms, seem to believe to be playing the first game, the one with unbounded credit, usually represented by the fortunes of the owners of the firms. In doing so, they have only been able to spot the seemingly “risk-less” profits lurking in the background.

In reality they have been playing the risky game with finite fortune $N < \infty$, unfortunately possessing no “free lunch”, and with a small, yet discernibly positive probability of a large loss. Such events sometimes materialize, at least according to theory, and history has confirmed that they also do in real life.

References

- [1] Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis. *Comm. Acad. Sci. Imp. Petropolitanae* 5. 175-192. (English translation: *Econometrica*, Vol 22, 23-36 (1954).
- [2] Buffon, G.L.L. (1777). Essai d'Arithmétique Morale. *Suppl. à l'Histoire Naturelle*, 46-148, Paris.
- [3] Feller, W. (1968). *An Introduction to Probability theory and its Applications*, Vol 1, 3rd ed. N.Y.

- [4] Martin-Löf, A. (1985). A Limit Theorem which classifies the “Petersburg Paradox”. *J. Appl. Prob.*, 22, 634-643.